

APPROXIMATE SCHUR COMPLEMENT PRECONDITIONING OF THE LOWEST-ORDER NODAL DISCRETIZATIONS*

J. DAVID MOULTON[†], JIM E. MOREL[‡], AND URI M. ASCHER[§]

Abstract. Certain classes of *nodal* methods and *mixed-hybrid* finite element methods lead to equivalent, robust, and accurate discretizations of second-order elliptic PDEs. However, widespread popularity of these discretizations has been hindered by the awkward linear systems which result. The present work overcomes this awkwardness and develops preconditioners which yield solution algorithms for these discretizations with an efficiency comparable to that of the multigrid method for standard discretizations.

Our approach exploits the natural partitioning of the linear system obtained by the mixed-hybrid finite element method. By eliminating different subsets of unknowns, two Schur complements are obtained with known structure. Replacing key matrices in this structure by lumped approximations, we define three optimal preconditioners. Central to the optimal performance of these preconditioners is their sparsity structure which is compatible with standard finite difference discretizations and hence treated adequately with only a single multigrid cycle.

In this paper we restrict the discussion to the two-dimensional case; these techniques are readily extended to three dimensions.

Key words. conjugate gradient, preconditioner, Schur complement, multigrid

AMS subject classifications. 65F10, 65N22, 65N30, 65N55

PII. S1064827596303491

1. Introduction. Nodal methods have long been one of the most popular discretization techniques employed within the reactor physics community to solve multi-group diffusion problems (see, e.g., [25, 13, 18, 24]). A survey of these methods can be found in [17]. Their success is a result of their exceptional accuracy which may be attributed to three distinct aspects of the *nodal ideology*. Specifically, akin to finite volume methods, nodal methods are physically motivated, cell-based discretizations that, by construction, rigorously enforce cell balance. They utilize an intriguing choice of unknowns (consisting, in two dimensions, for example, of cell and edge moments) which makes the nodal discretization naturally compatible with the various homogenization techniques that are crucial to reactor modeling. Moreover, comparable to *mixed finite element methods* (mixed FEM), the neutron current is obtained accurately and automatically from the discretization, thereby avoiding problematic finite difference approximations. Indeed, Hennart and del Valle [15] have shown that most nodal methods may be derived using *nonconforming* and *mixed-hybrid FEM* formulations, hence extending their realm of interest to the entire numerical analysis community.

*Received by the editors May 14, 1996; accepted for publication January 2, 1997. The U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes. Copyright is owned by SIAM to the extent not limited by these rights.

<http://www.siam.org/journals/sisc/19-1/30349.html>

[†]Department of Mathematics, University of British Columbia, Vancouver, British Columbia, Canada. Current address: Center for Nonlinear Studies, MS B258, Los Alamos, NM 87545 (moulton@cnls.lanl.gov). The research of this author was supported under NSERC Canada grant OGP0004306.

[‡]Scientific Computation, CIC-19, Los Alamos National Laboratory, Los Alamos, NM (jim@c3.lanl.gov). The research of this author was supported by U.S. Department of Energy contract W-7405-ENG-36.

[§]Department of Computer Science, University of British Columbia, Vancouver, British Columbia, Canada (ascher@cs.ubc.ca). The research of this author was partially supported under NSERC Canada grant OGP0004306.

However, widespread popularity of nodal discretizations has been hindered by the awkward linear systems which result. This is complicated further by the fact that each equivalent discretization (mixed, mixed-hybrid, and nonconformal FEMs) manifests this difficulty in a different way. For example, the mixed system is indefinite and although the elimination of certain unknowns leads to a positive definite system, the relative sparsity is compromised. Various techniques have been employed to circumvent this limitation (see, e.g., [23, 9, 1]). Conversely, the nonconformal FEMs are sparse but typically contain only edge unknowns, making the definition of a suitable grid hierarchy and the intergrid transfer operators difficult. Nevertheless, multilevel preconditioners have been developed for this system and include the multilevel substructuring preconditioners discussed in [10, 16]. However, the presence of strongly anisotropic diffusion in general may degrade the performance of these preconditioners. In this and most cases, solvers for the nonconformal discretization utilize the mixed-hybrid method only as a tool to establish equivalence with the corresponding mixed method. We contend that the mixed-hybrid method is in fact central to the development of efficient solvers as it is central to the equivalence of these discretizations.

The purpose of this work is to develop fast robust solvers for the nodal discretization (equivalently, the mixed-hybrid FEM) of linear self-adjoint second-order elliptic PDEs. To this end we will consider the various reduced systems mentioned previously, but more importantly we return to examine the underlying discretization in its indefinite form. Thus, we begin in section 2.1 by deriving a popular variant of the nodal discretizations, the constant-constant nodal integration method (NIM). The elimination of the currents and the cell-based unknowns is described along with the resulting edge-based system. We recall the mixed finite element methods in section 2.2 and, most importantly, their hybridization (i.e., the interelement continuity of the normal current $\mathbf{J} \cdot \mathbf{n}$ is treated as a constraint, imposed in the weak variational sense with Lagrange multipliers). Presentation of the mixed-hybrid FEM provides a transparent view of the indefinite system that is paramount in our work and obscured in the nodal derivation.

In section 3 we construct our preconditioners. We begin by eliminating the currents from the mixed-hybrid FEM system to obtain a reduced system for the scalar unknowns. Section 3.1 introduces the most common treatment of this reduced system, which eliminates the cell-based unknowns in favor of an exclusively edge-based system. This edge-based Schur complement is still sparse, and, moreover, is naturally partitioned by edge type (i.e., horizontal or vertical). However, a further reduction through the elimination of either edge type degrades the sparsity and hence a lumping procedure is introduced to approximate the Schur complement of this reduced system. The resulting approximation has a 9-point sparsity structure and is readily inverted with a standard multigrid solver such as black box multigrid [8]. We prove optimal convergence rates for this edge-based preconditioner. Unfortunately, these rates depend on the aspect ratio of the grid, and hence in section 3.2 we consider a *two-step* preconditioner which attempts to remove the directional bias of the lumping procedure. Ironically, this preconditioner is asymmetric, hence a more attractive possibility arises by returning to the reduced system for the scalar unknowns and proceeding to eliminate the edges (see section 3.3). Again a loss of sparsity results, and hence we apply a simple lumping to the original indefinite system. This lumping results in an approximate Schur complement which is the standard 5-point cell-centered discretization. Thus, multigrid inversion of this preconditioner is also possible. However, performing an *exact* multigrid solve for each iteration of preconditioned conjugate gra-

dient is undesirable, and hence in section 3.4 we discuss the effectiveness of a single multigrid cycle as an approximate inversion of the preconditioner.

Results of our numerical tests are reported in section 4 and include some nontrivial problems. In all cases, our preconditioners perform extremely well: only a small number of iterations, independent of the grid size, is required. The robustness of the *two-step* edge-based and the cell-based preconditioners is evident in the tests with high aspect ratio cells. Finally, conclusions are offered in section 5.

2. The discretizations. Since our primary interest is the development of preconditioners which exploit the underlying structure of the sparse linear systems, we restrict our discussion to linear second-order elliptic PDEs of the form (i.e., conservative one-group diffusion)

$$\begin{aligned} (2.1a) \quad & \nabla \cdot \mathbf{J} = Q(x, y), \\ (2.1b) \quad & \mathbf{J} = -D(x, y)\nabla\phi, \end{aligned}$$

where both the diffusion coefficient $D(x, y) > 0$ and the source $Q(x, y)$ are given functions on Ω , and the flux $\phi(x, y)$ is subject to the Dirichlet boundary condition

$$(2.2) \quad \phi = g(x, y) \quad \forall (x, y) \in \partial\Omega.$$

The discretizations which follow assume that Ω is a rectangular domain and employ an $L \times M$ tensor product mesh having cells $\Omega_{i,j} = (x_{i-\frac{1}{2}}, x_{i+\frac{1}{2}}) \times (y_{j-\frac{1}{2}}, y_{j+\frac{1}{2}})$. It is convenient to denote the mesh spacing by $\Delta x_i = x_{i+\frac{1}{2}} - x_{i-\frac{1}{2}}$ and $\Delta y_j = y_{j+\frac{1}{2}} - y_{j-\frac{1}{2}}$.

2.1. The nodal discretization. Common to all *nodal* discretizations is the choice of cell- and edge-based unknowns. Generally, these are taken to be moments up to some specified order, hence in the lowest-order case simple averages are employed. Specifically, the cell-based unknowns are averages defined by

$$(2.3) \quad \phi_{i,j} = \frac{1}{\Delta x_i \Delta y_j} \int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} \int_{y_{j-\frac{1}{2}}}^{y_{j+\frac{1}{2}}} \phi(x, y) dx dy$$

while the edge-based scalar unknowns, namely, edge averages of the scalar flux, are given by

$$(2.4) \quad \phi_{i+\frac{1}{2},j} = \phi_j(x_{i+\frac{1}{2}}), \quad \phi_j(x) = \frac{1}{\Delta y_j} \int_{y_{j-\frac{1}{2}}}^{y_{j+\frac{1}{2}}} \phi(x, y) dy$$

with the analogous definitions of $\phi_{i,j+\frac{1}{2}}$ and $\phi_i(y)$. Similarly, the edge-averaged currents are written as

$$(2.5) \quad J_{i+\frac{1}{2},j}^\pm = \lim_{x \rightarrow x_{i+\frac{1}{2}}^\pm} J_j(x), \quad J_j(x) = \frac{1}{\Delta y_j} \int_{y_{j-\frac{1}{2}}}^{y_{j+\frac{1}{2}}} \left[-D(x, y) \frac{\partial}{\partial x} \{ \phi(x, y) \} \right] dy$$

while $J_{i,j+\frac{1}{2}}^\pm$ is defined analogously.

All lowest-order members of the *nodal method* family, such as the nodal expansion method (NEM) [11], the nodal integration method (NIM) [12], the nodal Green's function method (NGFM) [19], and the coarse mesh expansion methods [20] yield equivalent discretizations of (2.1). We have chosen to present a brief discussion of the NIM. The first step in the NIM discretization consists of posing the cell-based

transverse integrated ODEs which govern $\phi_j(x)$ and $\phi_i(y)$. To this end we assume that an homogenized diffusion coefficient $D_{i,j}$ is defined on each cell, and transverse integrate (2.1) to obtain

$$-\frac{\partial}{\partial x} \left\{ D_{i,j} \frac{\partial}{\partial x} \phi_j(x) \right\} = -\frac{1}{\Delta y_j} \left\{ J_{j+\frac{1}{2}}^-(x) - J_{j-\frac{1}{2}}^+(x) \right\} + Q_j(x), \quad x_{i-\frac{1}{2}} \leq x \leq x_{i+\frac{1}{2}},$$

$$-\frac{\partial}{\partial y} \left\{ D_{i,j} \frac{\partial}{\partial y} \phi_i(y) \right\} = -\frac{1}{\Delta x_i} \left\{ J_{i+\frac{1}{2}}^-(y) - J_{i-\frac{1}{2}}^+(y) \right\} + Q_i(y), \quad y_{j-\frac{1}{2}} \leq y \leq y_{j+\frac{1}{2}},$$

where $J_{j\pm\frac{1}{2}}^\mp(x)$ and $J_{i\pm\frac{1}{2}}^\mp(y)$ are one-sided limits of the normal currents along the edges of the cell, and $Q_j(x)$ is defined in analogy with the transverse averaged unknowns. Thus, we have reduced the discretization of the PDE given in (2.1) to that of two ODEs that are coupled through *pseudo* source terms. The definition of the edge averages (2.4) naturally yields Dirichlet boundary conditions for each cell. Moreover, for the lowest-order or *constant-constant* NIM the *pseudo* source terms (i.e., $J_{j\mp\frac{1}{2}}^\pm(x)$ and $J_{i\mp\frac{1}{2}}^\pm(y)$) are assumed to be constant along their respective cell edges. By further assuming that the source $Q(x, y)$ is constant over each cell, we finally obtain a set of two *constant* coefficient ODEs with a constant *source* that are readily solved in terms of their values at the cell's edges (2.4).¹ With expressions for $\phi_j(x)$ and $\phi_i(y)$ in hand we construct the discretization. First, note that two independent definitions of the cell average are possible:

$$\phi_{i,j} = \frac{1}{\Delta x_i} \int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} \phi_j(x) dx = \frac{1}{\Delta y_j} \int_{y_{j-\frac{1}{2}}}^{y_{j+\frac{1}{2}}} \phi_i(y) dy$$

yielding *2LM* equations, e.g.,

$$\phi_{i,j} = \frac{1}{2}(\phi_{i+\frac{1}{2},j} + \phi_{i-\frac{1}{2},j}) - \frac{\Delta x_i^2}{12D_{i,j}} \left\{ \frac{1}{\Delta y_j} (J_{i,j+\frac{1}{2}}^- - J_{i,j-\frac{1}{2}}^+) - Q_{i,j} \right\}.$$

Furthermore, under the assumption of an homogenized diffusion coefficient and utilizing (2.1b) we obtain expressions for $J_{i\mp\frac{1}{2},j}^\pm, J_{i,j\mp\frac{1}{2}}^\pm$ on each cell, e.g.,

$$J_{i+\frac{1}{2},j}^- = -\frac{D_{i,j}}{\Delta x_i} (\phi_{i+\frac{1}{2},j} + \phi_{i-\frac{1}{2},j}) + \frac{1}{2\Delta x_i} \left\{ \frac{1}{\Delta y_j} (J_{i,j+\frac{1}{2}}^- - J_{i,j-\frac{1}{2}}^+) - Q_{i,j} \right\}.$$

Although these comprise four equations per cell, only three of them are linearly independent, as the same *balance* equation arises from both $J_{i+\frac{1}{2},j}^- - J_{i-\frac{1}{2},j}^+$ and $J_{i,j+\frac{1}{2}}^- - J_{i,j-\frac{1}{2}}^+$. Imposing continuity of $\mathbf{J} \cdot \mathbf{n}$ yields an equation for each interior edge (i.e., $(L-1)M + L(M-1)$ equations) while the boundary conditions give rise to $2L + 2M$ discrete boundary equations. Thus, we have $7LM + L + M$ equations in as many unknowns.

Although the *constant-constant* NIM discretization is complete at this point, it is seldom used in this form. Typically in the literature one proceeds by eliminating the edge currents $J_{i\mp\frac{1}{2},j}^\pm, J_{i,j\mp\frac{1}{2}}^\pm$, followed by the trivial elimination of the cell averages

¹If the PDE (2.1) includes the absorption term $\sigma\phi$, then the lower-order terms $\sigma_{i,j}\phi_j(x)$ and $\sigma_{i,j}\phi_i(y)$ are added accordingly to the above ODEs. Exact analytic solutions of the modified equations are still readily obtained.

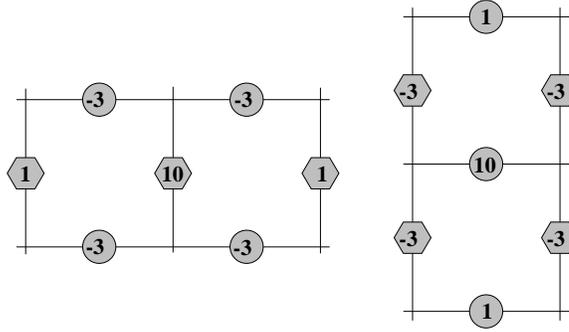


FIG. 2.1. The stencil weights for a uniform mesh with $D = 1$.

$\phi_{i,j}$ to obtain the 7-point nearest neighbor hexagonally coupled stencils that govern the edge-based fluxes $\phi_{i+\frac{1}{2},j}$:

$$\begin{aligned}
 (2.6) \quad & D_{i,j} \left(\frac{\Delta y_j}{\Delta x_i} \right) \left(\phi_{i+\frac{1}{2},j} - \phi_{i-\frac{1}{2},j} \right) - D_{i+1,j} \left(\frac{\Delta y_j}{\Delta x_{i+1}} \right) \left(\phi_{i+\frac{3}{2},j} - \phi_{i+\frac{1}{2},j} \right) \\
 & + \frac{3D_{i,j}\Delta x_i\Delta y_j}{\Delta x_i^2 + \Delta y_j^2} \left\{ \phi_{i+\frac{1}{2},j} + \phi_{i-\frac{1}{2},j} - \phi_{i,j+\frac{1}{2}} - \phi_{i,j-\frac{1}{2}} \right\} \\
 & + \frac{3D_{i+1,j}\Delta x_{i+1}\Delta y_j}{\Delta x_{i+1}^2 + \Delta y_j^2} \left\{ \phi_{i+\frac{3}{2},j} + \phi_{i+\frac{1}{2},j} - \phi_{i+1,j+\frac{1}{2}} - \phi_{i+1,j-\frac{1}{2}} \right\} \\
 & = \frac{1}{2} \frac{\Delta x_i\Delta y_j^3}{\Delta x_i^2 + \Delta y_j^2} Q_{i,j} + \frac{1}{2} \frac{\Delta x_{i+1}\Delta y_j^3}{\Delta x_{i+1}^2 + \Delta y_j^2} Q_{i+1,j}
 \end{aligned}$$

with the y -oriented rotated analogue at $\phi_{i,j+\frac{1}{2}}$ (see Figure 2.1).

2.2. Mixed-hybrid finite element methods. Recall that a typical *mixed* FEM views (2.1b) as a constraint coupled with the differential equation (2.1a), approximating a saddle point rather than a minimum [5]. This yields independent approximations of ϕ and \mathbf{J} . However, the resulting system is difficult to solve as it is indefinite, and, moreover, if it is recast in a positive definite form the relative sparsity is lost. *Hybrid* FEMs are a special class of *mixed* FEMs which temporarily relax certain interelement continuity conditions, ultimately enforcing them in only the weak variational sense as constraints. Of particular interest to this work is the interelement continuity of the normal current, $\mathbf{J} \cdot \mathbf{n}$, which may be treated as a constraint. Although hybridization of the mixed FEM also yields an indefinite system (it still approximates the saddle point), the distinct advantage is that it may be recast as a positive definite system in which the relative sparsity structure is preserved.

To be specific, we first define the function spaces for each unknown. Denote the set of edges E of the rectangles $\Omega_{i,j}$ as \mathcal{E}_h , so that defining $\mathcal{E}_h^\partial = \{E \in \mathcal{E}_h \mid E \subset \partial\Omega\}$ yields $\mathcal{E}_h^0 = \mathcal{E}_h \setminus \mathcal{E}_h^\partial$. Then the lowest-order Raviart–Thomas space may be written

$$\begin{aligned}
 RT_0^0(\Omega_h) &= \left\{ \mathbf{q} \mid \mathbf{q} \in RT_{-1}^0(\Omega_h), \int_{E^+} \mathbf{q} \cdot \mathbf{n} \, ds = \int_{E^-} \mathbf{q} \cdot \mathbf{n} \, ds \quad \forall E \in \mathcal{E}_h^0 \right\}, \\
 RT_{-1}^0(\Omega_h) &= \left\{ \mathbf{q} \mid \mathbf{q} \in L^2(\Omega) \times L^2(\Omega), \mathbf{q}|_{\Omega_{i,j}} \in \{ \{1, \xi\}, \{1, \eta\} \} \quad \forall \Omega_{i,j} \in \Omega_h \right\},
 \end{aligned}$$

where $RT_0^0 \subset H(\text{div}; \Omega)$ is the conforming subspace in which the solution \mathbf{J}_h will ultimately be found. The space of Lagrange multipliers $M_{-1}^0(\Omega_h)$ is defined as

$$M_{-1}^0(\Omega_h) = \{ \mu \mid \mu \in L^2(\Omega), \mu|_E \in \{1\} \quad \forall E \in \mathcal{E}_h^0; \mu|_E = 0 \quad \forall E \in \mathcal{E}_h^\partial \}$$

while the cell-based unknowns are defined over

$$U_{-1}^0(\Omega_h) = \left\{ \varphi \mid \varphi \in L^2(\Omega), \varphi|_{\Omega_{i,j}} \in \{1\} \quad \forall \Omega_{i,j} \in \Omega_h \right\}.$$

Thus, the discrete *mixed-hybrid* variational formulation of (2.1) may be expressed as *find* $(\mathbf{J}_h, \phi_h, \mu_h) \in RT_{-1}^0(\Omega_h) \times U_{-1}^0(\Omega_h) \times M_{-1}^0(\mathcal{E}_h^0)$ *such that*

$$\begin{aligned} \left(D^{-1} \mathbf{J}_h, \mathbf{q}_h \right) - \sum_{i,j} \left\{ (\phi_h, \nabla \cdot \mathbf{q}_h)_{\Omega_{i,j}} - \langle \mu_h, \mathbf{q}_h \cdot \mathbf{n} \rangle_{\partial \Omega_{i,j}} \right\} &= -\langle g, \mathbf{q}_h \cdot \mathbf{n} \rangle \quad \forall \mathbf{q}_h \in RT_{-1}^0(\Omega_h), \\ - \sum_{i,j} (\nabla \cdot \mathbf{J}_h, \varphi_h)_{\Omega_{i,j}} &= -(Q(x, y), \varphi_h) \quad \forall \varphi_h \in U_{-1}^0(\Omega_h), \\ \sum_{i,j} \langle \nu, \mathbf{J}_h \cdot \mathbf{n} \rangle_{\partial \Omega_{i,j}} &= 0 \quad \forall \nu_h \in M_{-2}^k(\mathcal{E}_h^0). \end{aligned}$$

Substituting the corresponding basis functions and integrating yields a system of the following form:

$$(2.7) \quad \begin{bmatrix} A & B^T & C^T \\ B & 0 & 0 \\ C & 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{J}_h \\ \phi_h \\ \boldsymbol{\mu}_h \end{bmatrix} = - \begin{bmatrix} \mathbf{Q}_J \\ \mathbf{Q}_\phi \\ 0 \end{bmatrix}.$$

Like the mixed FEM, this system is indefinite; however, unlike the mixed FEM, A is block diagonal with each block corresponding to a cell $\Omega_{i,j}$. Moreover, each block of A is itself block diagonal, with each block corresponding to a coordinate direction, hence A^{-1} retains the sparsity structure of A . Compared to *mixed* FEMs, a second-order global cell accuracy, rather than just at mid-cells, has been achieved, as it is possible to show that the Lagrange multipliers $\boldsymbol{\mu}_h$ are nothing more than the edge unknowns introduced in (2.4) and are second-order convergent [3]. Combining this with the cell average ϕ_h , which is also second-order convergent, it is possible to construct a *non-conforming* second-order approximation of $\phi(x, y)$ over each cell $\Omega_{i,j}$. Finally, it can be shown that upon elimination of the cell unknowns and the edge currents, the equations of the constant-constant NIM are obtained.

3. Reduced systems and approximate Schur complements. The linear system resulting from the nodal discretization (equivalently, the mixed-hybrid discretization) is indefinite. Fortunately this indefiniteness is readily circumvented with the block elimination of \mathbf{J}_h in (2.7) which yields

$$(3.1) \quad \mathcal{S}_{(\phi, \mu)} \begin{bmatrix} \phi_h \\ \boldsymbol{\mu}_h \end{bmatrix} = \mathbf{Q}_{\mathcal{S}_{(\phi, \mu)}}$$

with

$$\mathcal{S}_{(\phi, \mu)} = \begin{bmatrix} \mathcal{S}_B & BA^{-1}C^T \\ CA^{-1}B^T & \mathcal{S}_C \end{bmatrix}, \quad \mathbf{Q}_{\mathcal{S}_{(\phi, \mu)}} = - \begin{bmatrix} BA^{-1}\mathbf{Q}_J - \mathbf{Q}_\phi \\ CA^{-1}\mathbf{Q}_J \end{bmatrix},$$

where $\mathcal{S}_B = BA^{-1}B^T$ and $\mathcal{S}_C = CA^{-1}C^T$. Thus far the elementwise block diagonal structure of A and hence A^{-1} has preserved the relative sparsity of the system. The

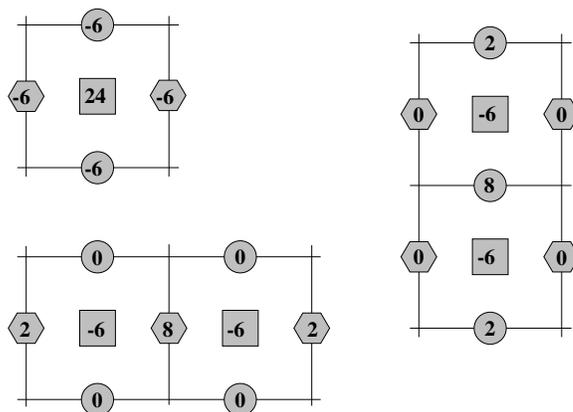


FIG. 3.1. The stencil weights of $\mathcal{S}_{(\phi, \mu)}$ on a uniform mesh with $D = 1$.

corresponding stencils are displayed in Figure 3.1, where the exact sparsity structure of $\mathcal{S}_{(\phi, \mu)}$ is apparent. Perhaps more importantly, the Schur complement $\mathcal{S}_{(\phi, \mu)}$ is symmetric positive definite, making it an ideal candidate for iterative methods. In addition, this reduced system may also be derived as a nonconformal finite element method by employing the canonical basis $\{1, \xi, \eta, \xi^2, \eta^2\}$. However, the majority of reported works using the nodal, mixed-hybrid and nonconformal methods proceed to eliminate the cell-based unknowns (see, e.g., [14, 16, 6]), so we consider this case first.

3.1. Lumping the edge-based Schur complement. Proceeding to eliminate the cell-based unknowns ϕ_h from (3.1) leads to the Schur complement of the edge unknowns

$$(3.2) \quad \mathcal{S}_\mu \boldsymbol{\mu}_h = \mathbf{Q}_{\mathcal{S}_\mu},$$

where

$$\begin{aligned} \mathcal{S}_\mu &= C(A^{-1} - A^{-1}B^T \mathcal{S}_B^{-1}BA^{-1})C^T, \\ \mathbf{Q}_{\mathcal{S}_\mu} &= C(A^{-1}B^T \mathcal{S}_B^{-1}BA^{-1} - A^{-1})\mathbf{Q}_J - CA^{-1}B^T \mathcal{S}_B^{-1}\mathbf{Q}_\phi. \end{aligned}$$

Note that (3.2) is precisely the edge-based formulation that was developed in section 2.1, given in (2.6), and shown schematically in Figure 2.1. The relative sparsity of the original system (2.7) continues to be preserved in \mathcal{S}_μ because \mathcal{S}_B is strictly diagonal (Figure 3.1). Yet, once again the most important property of this reduced system is that it is symmetric positive definite. In fact, it is this combination that has made (3.2) the most widely studied formulation of the mixed-hybrid finite element system (see, e.g., [16, 6]). However, as we shall see, it is not without its disadvantages.

To understand the inherent complexity of creating fast iterative solvers for (3.2) and to motivate our approach we partition the edge unknowns by orientation, vertical and horizontal, $\boldsymbol{\mu}_h = \{\mathbf{u}, \mathbf{v}\}$, such that (3.2) may be rewritten as

$$(3.3) \quad \begin{bmatrix} A_{uu} & A_{uv} \\ A_{vu} & A_{vv} \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} = \begin{bmatrix} \mathbf{Q}_u \\ \mathbf{Q}_v \end{bmatrix},$$

where A_{uu} , A_{vv} are tridiagonal and A_{uv} is bidiagonal ($A_{vu} = A_{uv}^T$). ADI methods appear to be a natural choice and have been investigated by a number of authors

[14, 4]. In addition, most standard preconditioners such as polynomial and incomplete factorizations have been studied [21]. However, none of these methods can attain the efficiency of multilevel solvers, the development of which has been hindered by the difficult hierarchy of grids and intergrid transfer operators which must be defined for these edge-based stencils (but see [2]). Certainly multilevel solvers would be much easier to develop if only a single orientation of the edge unknowns remained. However, the choice to eliminate either \mathbf{u} or \mathbf{v} is arbitrary, and moreover it is apparent that the respective Schur complements

$$\mathcal{S}_u = A_{uu} - A_{vu}(A_{vv})^{-1}A_{uv}, \quad \mathcal{S}_v = A_{vv} - A_{uv}(A_{uu})^{-1}A_{vu}$$

suffer an equivalent loss of sparsity as a result of the appearance of $(A_{vv})^{-1}$ and $(A_{uu})^{-1}$, respectively. Thus, posed purely as a problem in linear algebra, a suitable objective is to minimize the fill generated during the formation of \mathcal{S}_u (or \mathcal{S}_v). Further, noting that $(A_{uu})^{-1}$ is *pre-* and *post-*multiplied by bidiagonal matrices, the minimal fill attainable through approximations of $(A_{uu})^{-1}$ results when it is diagonal. Determining the *optimal* diagonal approximation is a foreboding task. However, recalling that the best incomplete Cholesky factorizations preserve *row sums*, it is reasonable to construct a diagonal approximation \widetilde{A}_{uu} which is composed simply of row sums of A_{uu} . A simple inversion then gives $(\widetilde{A}_{uu})^{-1} \approx (A_{uu})^{-1}$. Hence, the preconditioner may be written as

$$(3.4) \quad \widetilde{\mathcal{S}}_{\mu;u} = \begin{bmatrix} \widetilde{A}_{uu} & A_{uv} \\ A_{vu} & A_{vv} \end{bmatrix},$$

where we point out that the reduced system $\widetilde{\mathcal{S}}_v = A_{vv} - A_{vu}(\widetilde{A}_{uu})^{-1}A_{uv}$ is not only symmetric positive definite but also a second-order 9-point approximation of an elliptic PDE. Thus, efficient multigrid algorithms exist to “invert” it. In practice we use Dendy’s *black box* multigrid code [8], a robust solver which requires only the fine grid stencil as input.

3.1.1. A condition bound. We prove two results independently. The first is a proposition which makes the connection between the eigenvalues of the preconditioner and the difference between A_{uu} and its approximation.

PROPOSITION 3.1. *The eigenvalues of the preconditioned system $\lambda([\widetilde{\mathcal{S}}_{\mu;u}]^{-1}\mathcal{S}_\mu)$ may be written in the following form:*

$$\begin{aligned} \lambda &= 1, \text{ of multiplicity } N_v, \\ \lambda &= \lambda(\mathcal{S}_{\tilde{u}}^{-1}\mathcal{S}_u) = 1 + \lambda(\mathcal{S}_{\tilde{u}}^{-1}A_{uu}^*), \end{aligned}$$

where $\mathcal{S}_{\tilde{u}} = \widetilde{A}_{uu} - A_{uv}(A_{vv})^{-1}A_{vu}$ and $A_{uu}^* = A_{uu} - \widetilde{A}_{uu}$.

Proof. Observe that the inverse of the preconditioner $\widetilde{\mathcal{S}}_{\mu;u}$ may be expressed as

$$[\widetilde{\mathcal{S}}_{\mu;u}]^{-1} = \begin{bmatrix} \widetilde{A}_{uu} & A_{uv} \\ A_{vu} & A_{vv} \end{bmatrix}^{-1} = \begin{bmatrix} \mathcal{S}_{\tilde{u}}^{-1} & -\mathcal{S}_{\tilde{u}}^{-1}A_{uv}A_{vv}^{-1} \\ -[\widetilde{\mathcal{S}}_v]^{-1}A_{vu}[\widetilde{A}_{uu}]^{-1} & [\widetilde{\mathcal{S}}_v]^{-1} \end{bmatrix},$$

where $\widetilde{\mathcal{S}}_v = A_{vv} - A_{vu}(\widetilde{A}_{uu})^{-1}A_{uv}$. Block multiplication and simplification give

$$[\widetilde{\mathcal{S}}_{\mu;u}]^{-1}\mathcal{S}_\mu = \begin{bmatrix} \mathcal{S}_{\tilde{u}}^{-1}\mathcal{S}_u & 0 \\ -A_{vv}^{-1}A_{vu}\mathcal{S}_{\tilde{u}}^{-1}A_{uu}^* & I \end{bmatrix},$$

where $A_{uu}^* = A_{uu} - \widetilde{A_{uu}}$. Hence, the eigenvalues are $\lambda = 1$ with multiplicity N_v and $\lambda = \lambda(\mathcal{S}_{\tilde{u}}^{-1}\mathcal{S}_u)$. Since $\mathcal{S}_u = \mathcal{S}_{\tilde{u}} + A_{uu}^*$ we have $\mathcal{S}_{\tilde{u}}^{-1}\mathcal{S}_\mu = I + \mathcal{S}_{\tilde{u}}^{-1}A_{uu}^*$. \square

Next, we bound the condition number of the edge-based preconditioned system, independently of the mesh size and the diffusion coefficient.

THEOREM 3.2. *Let $r_{i,j} = \Delta y_j / \Delta x_i$. Then the condition number of the preconditioned system is bounded by*

$$(3.5) \quad \kappa\left([\widetilde{\mathcal{S}_{\mu;u}}]^{-1}\mathcal{S}_\mu\right) \leq \max_{r_{i,j}} \left\{ (1 + r_{i,j}^2), \frac{3}{(1 + r_{i,j}^2)} \right\}.$$

With a uniform grid spacing $r = r_{i,j}$, an even tighter bound is realized:

$$(3.6) \quad \kappa\left([\widetilde{\mathcal{S}_{\mu;u}}]^{-1}\mathcal{S}_\mu\right) \leq \max \left\{ \frac{1}{3}(1 + r^2), \frac{3}{(1 + r^2)} \right\}.$$

Proof. We employ a *superelement analysis* as proposed by Kuznetsov and Maliassov [16] to obtain a local bound on the condition number κ . Consider first the elemental construction of the edge-based system by either utilizing its equivalence with a nonconformal method or by elementwise elimination of the currents and the cell-based unknowns from the original system (2.7). On a single cell $\Omega_{i,j}$ we have

$$\mathcal{S}_\mu^s = \begin{bmatrix} A_{uu}^s & A_{uv}^s \\ A_{vu}^s & A_{vv}^s \end{bmatrix} = \left[\begin{array}{cc|cc} \beta + \alpha & \beta - \alpha & -\beta & -\beta \\ \beta - \alpha & \beta + \alpha & -\beta & -\beta \\ \hline -\beta & -\beta & \beta + \gamma & \beta - \gamma \\ -\beta & -\beta & \beta - \gamma & \beta + \gamma \end{array} \right]$$

with $\alpha = D_{i,j}\Delta y_j / \Delta x_i$, $\beta = 3D_{i,j}\Delta x_i\Delta y_j / (\Delta x_i^2 + \Delta y_j^2)$, $\gamma = D_{i,j}\Delta x_i / \Delta y_j$, and where the unknowns on $\Omega_{i,j}$ have been ordered as $[\phi_{i-\frac{1}{2},j}, \phi_{i+\frac{1}{2},j}, \phi_{i,j-\frac{1}{2}}, \phi_{i,j+\frac{1}{2}}]$. The lumped approximation gives

$$[\widetilde{\mathcal{S}_{\mu;u}}]^s = \begin{bmatrix} [\widetilde{A_{uu}}]^s & A_{uv}^s \\ A_{vu}^s & A_{vv}^s \end{bmatrix}, \quad [\widetilde{A_{uu}}]^s = \begin{bmatrix} 2\beta & 0 \\ 0 & 2\beta \end{bmatrix}.$$

It is straightforward to verify $\mathcal{K} = \ker(\mathcal{S}_\mu^s) = \ker([\widetilde{\mathcal{S}_{\mu;u}}]^s) = [1, 1, 1, 1]^T$. This kernel is to be expected as \mathcal{S}_μ^s essentially defines a pure Neumann problem on the cell $\Omega_{i,j}$. Finally, we consider the generalized eigenvalue problem

$$\mathcal{S}_{\mu_s}^s \boldsymbol{\mu}_s = \lambda_s [\widetilde{\mathcal{S}_{\mu;u}}]^s \boldsymbol{\mu}_s \quad \text{with } \boldsymbol{\mu}_s \perp \mathcal{K}.$$

Applying Theorem 6.3.1 of Rao and Mitra [22] we find that the three ‘‘proper’’ eigenvalues are $\lambda_s = [1, 1, \frac{1}{3}(1 + r_{i,j}^2)]$. If $r_{i,j} > \sqrt{2}$, then $\lambda_s \in [1, \frac{1}{3}(1 + r_{i,j}^2)]$, while if $r_{i,j} < \sqrt{2}$, then $\lambda_s \in [\frac{1}{3}(1 + r_{i,j}^2), 1]$. Hence, if the mesh spacing is constant, then *superelement analysis* trivially yields the bound given in (3.6). Conversely, on a spatially dependent mesh we consider the following:

$$\begin{aligned} r_{i,j} > \sqrt{2} \quad \forall(i,j) &\Rightarrow \kappa \leq \max_{r_{i,j}} \left\{ \frac{1}{3}(1 + r_{i,j}^2) \right\}, \\ r_{i,j} < \sqrt{2} \quad \forall(i,j) &\Rightarrow \kappa \leq \max_{r_{i,j}} \left\{ \frac{3}{(1 + r_{i,j}^2)} \right\}, \\ r_{i,j} > \sqrt{2} \text{ and } r_{k,l} < \sqrt{2} &\Rightarrow \kappa \leq \max_{r_{i,j}, r_{k,l}} \left\{ \frac{1 + r_{i,j}^2}{1 + r_{k,l}^2} \right\} \leq \max_{r_{i,j}} \{1 + r_{i,j}^2\}. \end{aligned}$$

Combining these bounds yields (3.5). \square

The form of the bound given in (3.5) is unchanged when the analysis is extended to a diagonal diffusion tensor. The crucial difference is the definition of $r_{i,j}$, which is modified to read

$$r_{i,j} = \sqrt{\frac{D_{i,j}^{(x)}}{D_{i,j}^{(y)}} \left[\frac{\Delta y_j}{\Delta x_i} \right]}.$$

Hence, we see that the implication of poor conditioning on a spatially dependent high aspect ratio grid is also indicative of poor conditioning in the presence of strong anisotropy.

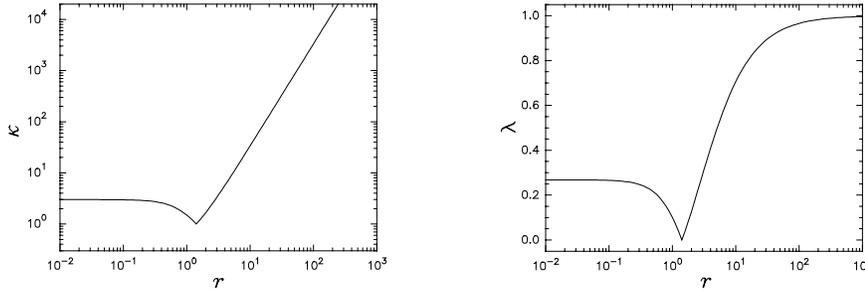


FIG. 3.2. Upper bound of the condition number, κ (equation (3.6)), and of the decay rate $\lambda_\infty(r)$.

3.1.2. Bounding the elements of A_{uu}^{-1} . The utility of searching for diagonal approximations of A_{uu} which eliminate the dependence on $r_{i,j}$ may be understood if we employ Theorem 2.4 of Demko, Moss, and Smith [7] which bounds the exponential decay rate of elements in the inverse of a banded matrix. Specifically, we consider constant mesh spacing and constant diffusivity with Dirichlet boundary conditions such that the diagonal blocks of A_{uu} are given by

$$(3.7) \quad (A_{uu})_{jj} = D \left(\frac{r}{1+r^2} \right) \text{tri} [2-r^2, 8+2r^2, 2-r^2], \quad j = 1, \dots, M,$$

and are of dimension $(L-1) \times (L-1)$. A bound on the elemental decay rate may be written in the form

$$(3.8) \quad \left| \left\{ (A_{uu})_{jj}^{-1} \right\}_{kl} \right| \leq C \lambda(r)^{|k-l|},$$

where the decay rate is itself bounded

$$(3.9) \quad \lambda(r) \leq \lambda_\infty(r) = \frac{\alpha^+ - \alpha^-}{\alpha^+ + \alpha^-}, \quad \alpha^\pm = \sqrt{(4+r^2) \pm |2-r^2|}.$$

We note that the bound λ_∞ is consistent with Theorem 3.2, clearly capturing its dependence on r . In particular, the exact nature of the approximation for $r = \sqrt{2}$ is observed with $\lambda(\sqrt{2}) = 0$. But perhaps most importantly, Figure 3.2 reveals that $\lambda_\infty(r) \rightarrow 1$ as $r \rightarrow \infty$, implying that the diagonal blocks of A_{uu}^{-1} are not only dense, but all the elements are of approximately the same magnitude. Thus, it is not surprising that a diagonal approximation is inadequate.

3.2. A two-step lumped approximation. Eliminating \mathbf{u} to obtain the Schur complement for \mathbf{v} is an arbitrary choice and it may not always be the natural one. Given this preconditioner's dependence on the ordering of unknowns, we postulate that it can be improved, especially for nonuniform grids, in much the same manner that symmetric SOR (SSOR) improves upon SOR (i.e., by composing the two possible orderings). To this end consider two splittings of \mathcal{S}_μ ,

$$\mathcal{S}_\mu = E_1 + F_1, \quad \mathcal{S}_\mu = E_2 + F_2,$$

which are employed in the following two-stage iteration of the system $\mathcal{S}_\mu \boldsymbol{\mu}_h = \mathbf{Q}_{\mathcal{S}_\mu}$:

$$(3.10a) \quad E_1 \boldsymbol{\mu}_h^{(k+\frac{1}{2})} = -F_1 \boldsymbol{\mu}_h^{(k)} + \mathbf{Q}_{\mathcal{S}_\mu},$$

$$(3.10b) \quad E_2 \boldsymbol{\mu}_h^{(k+1)} = -F_2 \boldsymbol{\mu}_h^{(k+\frac{1}{2})} + \mathbf{Q}_{\mathcal{S}_\mu}.$$

Defining the residual in the standard manner, $\mathbf{r}^{(k)} = \mathbf{Q}_{\mathcal{S}_\mu} - \mathcal{S}_\mu \boldsymbol{\mu}_h^{(k)}$, allows (3.10a) to be rewritten in the form

$$(3.11) \quad \boldsymbol{\mu}_h^{(k+\frac{1}{2})} = \boldsymbol{\mu}_h^{(k)} + E_1^{-1} \mathbf{r}^{(k)}$$

such that substitution of (3.11) into (3.10b) gives

$$E_2 \boldsymbol{\mu}_h^{(k+1)} = -F_2 \left(\boldsymbol{\mu}_h^{(k)} + E_1^{-1} \mathbf{r}^{(k)} \right) + \mathbf{Q}_{\mathcal{S}_\mu} = E_2 \boldsymbol{\mu}_h^{(k)} + (E_1 - F_2) E_1^{-1} \mathbf{r}^{(k)}.$$

Hence, the two-stage iteration is readily expressed as a single residual equation

$$(3.12) \quad \left\{ E_1 (E_1 - F_2)^{-1} E_2 \right\} \Delta \boldsymbol{\mu}_h^{(k+1)} = \mathbf{r}^{(k)}.$$

To apply this to the *nodal* equations we need only realize that a splitting exists for the *lumped* preconditioner. The splitting is described by

$$E_1 = \begin{bmatrix} \widetilde{A}_{uu} & A_{uv} \\ A_{vu} & A_{vv} \end{bmatrix}, \quad F_1 = \begin{bmatrix} A_{uu}^* & 0 \\ 0 & 0 \end{bmatrix}$$

with $A_{uu} = \widetilde{A}_{uu} + A_{uu}^*$. Similarly, we introduce the lumping in \mathbf{v} writing $A_{vv} = \widetilde{A}_{vv} + A_{vv}^*$. Substitution into (3.12) yields a *two-step lumped* preconditioner

$$(3.13) \quad \begin{aligned} \widetilde{\mathcal{S}}_{\mu;u,v} &= \left\{ \widetilde{\mathcal{S}}_{\mu;u} [\widetilde{\mathcal{S}}_\mu]^{-1} \widetilde{\mathcal{S}}_{\mu;v} \right\} \\ &= \begin{bmatrix} \widetilde{A}_{uu} & A_{uv} \\ A_{vu} & A_{vv} \end{bmatrix} \begin{bmatrix} \widetilde{A}_{uu} & A_{uv} \\ A_{vu} & \widetilde{A}_{vv} \end{bmatrix}^{-1} \begin{bmatrix} A_{uu} & A_{uv} \\ A_{vu} & A_{vv} \end{bmatrix}. \end{aligned}$$

Observe that $\widetilde{\mathcal{S}}_{\mu;v} \neq \widetilde{\mathcal{S}}_{\mu;u}^T$ and hence $\widetilde{\mathcal{S}}_{\mu;u,v}$ will be asymmetric in general. This is unfortunate, as it precludes the use of preconditioned conjugate gradients used elsewhere in this work. However, there are many alternative, transform-free Krylov subspace solvers available. In our numerical tests we employ preconditioned ORTHORES [26] with a truncation length of only five and observe excellent performance.

The following result again relates the eigenvalues of the preconditioned system to the approximation difference in the Schur complement and, like Proposition 3.1, can be easily proved.

PROPOSITION 3.3. *The eigenvalues of the preconditioned system arising from the application of $\widetilde{\mathcal{S}}_{\mu;u,v}$ are given by*

$$\begin{aligned} \lambda &= 1, \text{ of multiplicity } N_v, \\ \lambda &= \lambda \left(\left[\widetilde{\mathcal{S}}_{\tilde{u}}^{-1} \widetilde{\mathcal{S}}_u \right]^{-1} \left[\mathcal{S}_{\tilde{u}}^{-1} \mathcal{S}_u \right] \right), \\ &= 1 - \lambda \left(\left[\widetilde{\mathcal{S}}_u^{-1} A_{uv} (\widetilde{A}_{vv})^{-1} A_{vv}^* \right] \left[\widetilde{\mathcal{S}}_v^{-1} A_{vu} (\widetilde{A}_{uu})^{-1} A_{uu}^* \right] \right), \end{aligned}$$

where $\widetilde{\mathcal{S}}_{\tilde{u}} = \widetilde{A}_{uu} - A_{uv} (\widetilde{A}_{vv})^{-1} A_{vu}$.

Unfortunately, in general, superelement analysis is not directly applicable to this preconditioner. However, it is very encouraging that the eigenvalues may be expressed as a perturbation from unity, with the perturbation a product involving the errors associated with the lumped approximations. These errors possess a cell-based block diagonal structure; hence for every cell at least one of A_{uu}^* and A_{vv}^* is associated with a good approximation. Thus, we anticipate $\widetilde{\mathcal{S}}_{\mu;u,v}$ to be robust for high aspect ratio problems, as is demonstrated computationally in section 4.

3.3. Lumping the cell-based Schur complement. Eliminating $\boldsymbol{\mu}_h$ from (3.1) we obtain the Schur complement for the cell-based unknowns

$$(3.14) \quad \mathcal{S}_\phi \boldsymbol{\phi}_h = \mathbf{Q}_{\mathcal{S}_\phi},$$

where

$$\begin{aligned} \mathcal{S}_\phi &= B(A^{-1} - A^{-1}C^T \mathcal{S}_C^{-1}CA^{-1})B^T, \\ \mathbf{Q}_{\mathcal{S}_\phi} &= \mathbf{Q}_\phi - B(A^{-1} - A^{-1}C^T \mathcal{S}_C^{-1}CA^{-1})\mathbf{Q}_J. \end{aligned}$$

The Schur complement \mathcal{S}_ϕ is symmetric positive definite and hence we consider the development of a preconditioner for it. Of particular importance to the robustness of the resulting solvers is that this reduced system (3.14) governs only the cell-based unknowns. However, the relative sparsity structure of the system has been compromised because \mathcal{S}_C is block diagonal, with each block itself tridiagonal. Thus, the matrix-vector multiplication $\mathcal{S}_\phi \boldsymbol{\phi}_h$ will involve the inversion of M , $(L-1) \times (L-1)$ and L , $(M-1) \times (M-1)$ tridiagonal matrices (see the remark below).

To develop a preconditioner for \mathcal{S}_ϕ and, in particular, one which is suitable for multigrid inversion, we return to the full system (2.7) and note that $\mathcal{S}_\phi > 0$ if $A > 0$ and $\text{range}(B) \cap \text{range}(C) = \{\emptyset\}$. Hence, the primary objective is the minimization of fill through symmetric positive definite approximations of A . To this end we continue the idea of the previous section and replace A , a block diagonal matrix, with \tilde{A} , a strictly diagonal lumped approximation:

$$A_{i,j} = \left\{ \frac{\Delta x_i \Delta y_i}{6D_{i,j}} \right\} \begin{bmatrix} 2 & 1 & 0 & 0 \\ 1 & 2 & 0 & 0 \\ 0 & 0 & 2 & 1 \\ 0 & 0 & 1 & 2 \end{bmatrix} \longrightarrow \tilde{A}_{i,j} = \left\{ \frac{\Delta x_i \Delta y_i}{2D_{i,j}} \right\} I_4,$$

where I_4 is the 4×4 identity matrix. This simple approximation transforms $\mathcal{S}_\phi \rightarrow \widetilde{\mathcal{S}}_\phi$, where $\widetilde{\mathcal{S}}_\phi$ is in fact the well-known 5-point cell-centered discretization of (2.1). Thus, we have derived an approximate Schur complement which may be used as a preconditioner and which may again be efficiently inverted using a standard multigrid method.

3.3.1. A condition bound. The use of tridiagonal solves in the matrix–vector product $\mathcal{S}_\phi \phi_h$ makes a cell-based superelement analysis impossible. Certainly using superelements defined over entire rows and columns of the mesh is possible, but obtaining analytic expressions for the corresponding eigenvalues seems improbable. It is in fact more interesting to consider the full scalar system (3.1) for which we bound the condition number of the correspondingly preconditioned system by a constant independent of the mesh and the diffusion coefficient.

THEOREM 3.4. *The preconditioner $\widetilde{\mathcal{S}}_{(\phi,\mu)}$ yields an effective condition number which is bounded by*

$$(3.15) \quad \kappa \left(\left[\widetilde{\mathcal{S}}_{(\phi,\mu)} \right]^{-1} \mathcal{S}_{(\phi,\mu)} \right) \leq 3$$

independently of the mesh and the diffusion coefficient.

The bound (3.15) is unchanged if we extend the analysis to a diagonal diffusion tensor. The implication is that this preconditioner is robust with respect to anisotropy as well as spatial variations in the diffusion tensor.

Proof. Once again we employ *superelement analysis* to relate local bounds, which are readily obtained, to global ones. We begin by writing down the elemental problem

$$\mathcal{S}_{(\phi,\mu)}^s = \left[\begin{array}{cc|cc|cc} 12(\alpha + \gamma) & & -6\alpha & -6\alpha & -6\gamma & -6\gamma \\ -6\alpha & & 4\alpha & 2\alpha & 0 & 0 \\ -6\alpha & & 2\alpha & 4\alpha & 0 & 0 \\ \hline -6\gamma & & 0 & 0 & 4\gamma & 2\gamma \\ -6\gamma & & 0 & 0 & 2\gamma & 4\gamma \end{array} \right]$$

with $\alpha = D_{i,j} \Delta y_j / \Delta x_i$, $\gamma = D_{i,j} \Delta x_i / \Delta y_j$, and where the unknowns have been ordered $\Phi_s = [\phi_{i,j}, \phi_{i-\frac{1}{2},j}, \phi_{i+\frac{1}{2},j}, \phi_{i,j-\frac{1}{2}}, \phi_{i,j+\frac{1}{2}}]^T$. Approximating A by \tilde{A} gives

$$\tilde{\mathcal{S}}_{(\phi,\mu)}^s = \left[\begin{array}{cc|cc|cc} 4(\alpha + \gamma) & & -2\alpha & -2\alpha & -2\gamma & -2\gamma \\ -2\alpha & & 2\alpha & 0 & 0 & 0 \\ -2\alpha & & 0 & 2\alpha & 0 & 0 \\ \hline -2\gamma & & 0 & 0 & 2\gamma & 0 \\ -2\gamma & & 0 & 0 & 0 & 2\gamma \end{array} \right],$$

where it is apparent that we have preserved the kernel

$$\mathcal{K} = \ker(\mathcal{S}_{(\phi,\mu)}^s) = \ker(\tilde{\mathcal{S}}_{(\phi,\mu)}^s) = [1, 1, 1, 1, 1]^T.$$

Solving the corresponding generalized eigenvalue problem

$$\mathcal{S}_{(\phi,\mu)}^s \Phi_s = \lambda_s \tilde{\mathcal{S}}_{(\phi,\mu)}^s \Phi_s \quad \text{with } \Phi_s \perp \mathcal{K}$$

we obtain $\lambda_s = [1, 1, 3, 3]$. Since these eigenvalues are independent of the cell indices (i, j) the global bound follows immediately. \square

Remark. Based on the equivalence of mixed and mixed-hybrid FEMs it is apparent that \mathcal{S}_ϕ is equivalently the reduced system obtained in mixed FEMs. Generally, it is problematic to solve this system because of the loss of sparsity. Here we are able to take advantage of the restricted geometry, incurring only the additional cost of tridiagonal solves. A sensitivity to highly variable coefficients has also been a problem for some methods. Allan, Ewing, and Lu [1] observed this and employed a lumped

diagonal preconditioner to the mixed formulation, and in taking advantage of the assumed rectangular geometry, obtain results comparable to ours. However, with the use of the mixed-hybrid discretization we have found a simpler approach to a robust cell-based preconditioner which in the form $\widetilde{\mathcal{S}}_{(\phi,\mu)}$ has the potential to be extended to more general geometries.

3.4. Approximate multilevel inversion of the preconditioner. Using either of the aforementioned preconditioners would be clearly impractical if we actually intended to “invert” them to some strict level of accuracy. However, this is unnecessary because the preconditioners are standard discretizations of (2.1) and hence standard multigrid methods yield an error reduction per iteration which is independent of the grid size. This implies a spectral equivalence which allows us to employ a single V- or W-cycle in lieu of a complete multigrid solve. We state and prove the following theorem for completeness.

THEOREM 3.5. *Let $\mathcal{A}, \mathcal{B}, \mathcal{C}$ be $n \times n$ symmetric positive definite matrices, and assume*

$$(3.16) \quad \rho(I - \mathcal{C}^{-1}\mathcal{B}) \leq c < 1,$$

where ρ denotes the spectral radius. Then

$$(3.17) \quad \kappa(\mathcal{C}^{-1}\mathcal{A}) \leq \left\{ \frac{1+c}{1-c} \right\} \kappa(\mathcal{B}^{-1}\mathcal{A}).$$

Proof. Let λ be an eigenvalue of $\mathcal{C}^{-1}\mathcal{B}$. Then (3.16) implies $1 - c \leq \lambda \leq 1 + c$ and hence we readily obtain (3.17). \square

The practical implication of Theorem 3.5 follows if we let \mathcal{A} be the edge-based \mathcal{S}_μ or the cell-based \mathcal{S}_ϕ , and let \mathcal{B} be its symmetric preconditioner. \mathcal{B} is further approximated by \mathcal{C} , a single multigrid cycle for which we have c independent of grid size, typically $c \approx 0.1$. Thus, $\frac{1+c}{1-c} \leq 1.3$, implying that \mathcal{C} is (almost) as effective as \mathcal{B} .

4. Numerical tests. A progressive test suite is presented which systematically highlights the strengths and weaknesses of each approximate Schur complement preconditioner. Beginning with a constant coefficient Dirichlet problem on the unit square (section 4.1) we verify numerically that on a uniform mesh ($r = 1$) all three preconditioners exhibit mesh independent convergence. Continuing with this simple example we vary the aspect ratio of a constant mesh to confirm the breakdown of the simple edge-based preconditioner $\widetilde{\mathcal{S}}_{\mu;u}$ and the robustness of $\widetilde{\mathcal{S}}_{\mu;u,v}$ and $\widetilde{\mathcal{S}}_\phi$. In the next test we consider a ground water flow problem (section 4.2) with significant jumps in the diffusion coefficient. To isolate the influence of this spatial dependence we solve this problem on a uniform grid, once again observing mesh independent convergence. To evaluate the relative cost of each preconditioner machine timings are also presented. Finally, in section 4.3 we present a diffusive checkerboard problem which combines significant jumps in $D(x, y)$ with a spatially dependent grid containing high aspect ratio cells. The inevitable breakdown of $\widetilde{\mathcal{S}}_{\mu;u}$ is observed along with the robustness of $\widetilde{\mathcal{S}}_{\mu;u,v}$ and $\widetilde{\mathcal{S}}_\phi$.

4.1. A toy problem. Consider the Dirichlet problem with constant diffusivity $D(x, y) \equiv 1$ and solution

$$\phi(x, y) = 2 + \sin(\alpha\pi x) \sin(\beta\pi y),$$

TABLE 4.1

A comparison of iteration counts for the two preconditioners, $\widetilde{\mathcal{S}}_{\mu;u}$ and $\widetilde{\mathcal{S}}_{\phi}$ for the toy problem with constant mesh spacing on the domain $[0, 1] \times [0, 1]$. (Convergence criteria of 10^{-6} .)

Mesh size	$\widetilde{\mathcal{S}}_{\mu;u}$		$\widetilde{\mathcal{S}}_{\mu;u,v}$		$\widetilde{\mathcal{S}}_{\phi}$	
	C	V	C	V	C	V
20×20	6	6	4	4	11	11
40×40	6	6	4	4	11	11
80×80	6	6	4	4	11	11

TABLE 4.2

Iteration counts for the lumped preconditioner, $\widetilde{\mathcal{S}}_{\mu;u}$ with constant mesh spacing on the domain $[0, a] \times [0, b]$ and a convergence criteria of 10^{-6} . (Note: $\alpha = 2/a, \beta = 2/b$.)

Mesh size	Aspect ratio: $r = b/a$						
	1/8	1/4	1/2	1	2	4	8
10×10	9	9	9	6	7	15	19
20×20	9	10	9	6	7	16	30
40×40	9	10	9	6	7	15	29

where $\{\alpha, \beta\}$ are free parameters. Substitution of the solution into (2.1) yields the required source

$$Q(x, y) = (\alpha^2 + \beta^2) \pi^2 \sin(\alpha\pi x) \sin(\beta\pi y).$$

A comparison of the two edge-based preconditioners and the cell-based preconditioner is presented in Table 4.1 for the domain $[0, 1] \times [0, 1]$ with $\alpha = \beta = 2$. It is clear from these results that all preconditioners generate an average residual reduction which is independent of the mesh size. Moreover, there is no difference in the iteration counts for those runs which inverted the preconditioner *exactly* (i.e., columns marked C) and those which only used a single V(1,1,1)-cycle (i.e., columns marked V). Thus, these solvers offer an efficiency for the solution of the nodal equations which is comparable to standard multigrid algorithms applied to standard discretizations of Poisson equations. Moreover, these results are consistent with the theoretical analysis performed in section 3.

Unfortunately, we must also demonstrate the vulnerability of $\widetilde{\mathcal{S}}_{\mu;u}$ regarding high aspect ratio cells, and hence we conducted several runs in which the size of the physical domain ($[0, a] \times [0, b]$) was varied while the mesh size remained fixed. The results for the preconditioned edge-based solver $\widetilde{\mathcal{S}}_{\mu;u}$ are summarized in Table 4.2. Here it is apparent that for r sufficiently close to one, this solver demonstrates excellent efficiency. Performance is still excellent as r approaches zero. However, just as the bound on κ (Theorem 3.2, equation (3.6)) predicts, the efficiency is degraded as r increases. Conversely, the two-step lumped preconditioner $\widetilde{\mathcal{S}}_{\mu;u,v}$ displays perfectly symmetric iteration counts on the 40×40 mesh (Table 4.3). This symmetry in r is very encouraging, particularly in light of the preconditioners asymmetry. Similarly, Table 4.4 clearly displays the r -independent convergence of $\widetilde{\mathcal{S}}_{\phi}$.

4.2. Ground water flow. The saturated flow problem of Mosé et al. [21] serves both as an excellent test of the preconditioners developed in section 3 and as a showcase for the discretizations themselves. The problem, shown in Figure 4.1, models the

TABLE 4.3

Iteration counts for the two-step lumped preconditioner, $\widetilde{\mathcal{S}}_{\mu;u,v}$ with constant mesh spacing on the domain $[0, a] \times [0, b]$ and a convergence criteria of 10^{-6} . (Note: $\alpha = 2/a, \beta = 2/b$.)

Mesh size	Aspect ratio: $r = b/a$						
	1/8	1/4	1/2	1	2	4	8
10×10	7	7	5	4	5	6	7
20×20	9	7	5	4	5	6	8
40×40	9	7	5	4	5	7	9

TABLE 4.4

Iteration counts for the cell-based preconditioner, $\widetilde{\mathcal{S}}_{\phi}$ with constant mesh spacing on the domain $[0, a] \times [0, b]$ and a convergence criteria of 10^{-6} . (Note: $\alpha = 2/a, \beta = 2/b$.)

Mesh size	Aspect ratio: $r = b/a$						
	1/8	1/4	1/2	1	2	4	8
10×10	9	10	10	10	10	10	9
20×20	10	10	11	11	11	11	10
40×40	11	11	11	11	11	11	11

flow of a ground water through the channel formed by the impervious sides,

$$\mathbf{J} \cdot \mathbf{n} = 0 \quad \forall (x, y) \in \Gamma_I = \{x = 0, x = 100, 0 \leq y \leq 100\}$$

and driven by the externally imposed gradient

$$(4.1) \quad \phi(x, y) = 100 \quad \forall (x, y) \in \Gamma_T = \{0 \leq x \leq 100, y = 100\},$$

$$(4.2) \quad \phi(x, y) = 99 \quad \forall (x, y) \in \Gamma_B = \{0 \leq x \leq 100, y = 0\}.$$

The flow is impeded by two regions of extremely low hydraulic conductivity. Solutions computed on a very coarse mesh of 25×25 are shown in Figure 4.2. Particularly impressive are the streamlines which are significantly more accurate than those obtained even with careful postprocessing of conforming methods [21]. Computations were performed on a uniform mesh with the resulting iteration counts presented in Table 4.5. The columns marked $CG(\mathcal{S}_{\mu})$ and $CG(\mathcal{S}_{\phi})$ record the performance of conjugate gradients applied directly to the reduced systems \mathcal{S}_{μ} and \mathcal{S}_{ϕ} , respectively. This provides an indication of the problem's conditioning. For the purpose of comparison, iteration counts for diagonal preconditioning of the edge-based system, denoted \mathcal{S}_{μ}^D , are also included. Consistent with the theoretical bounds, the entries for the three approximate Schur complement preconditioners are identical to the constant coefficient Dirichlet problem given in Table 4.1.

In Table 4.6 we present machine timings for the three preconditioners². Not surprisingly, the simplest preconditioner $\widetilde{\mathcal{S}}_{\mu;u}$, which in this case also yields the smallest bound on the condition of the system, is the fastest. However, we have established the lack of robustness in $\widetilde{\mathcal{S}}_{\mu;u}$ and it is encouraging to see that only approximately 50% more time is required by the more robust preconditioners $\widetilde{\mathcal{S}}_{\mu;u,v}$ and $\widetilde{\mathcal{S}}_{\phi}$. It is also interesting to note that despite having approximately twice as many unknowns in the iterative process *and* needing OrthoRes in place of a simple conjugate gradient solver, $\widetilde{\mathcal{S}}_{\mu;u,v}$ is competitive with $\widetilde{\mathcal{S}}_{\phi}$. This is of particular interest to people with existing

²Timings were performed on an HP Apollo 9000/735 (99 MHz PA-RISC 7100).

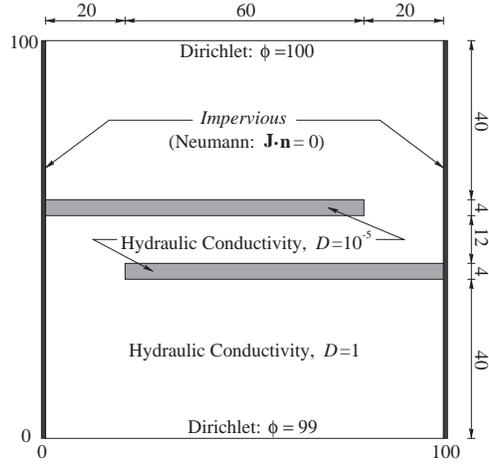


FIG. 4.1. A schematic of the first example in Mosé et al. [21].

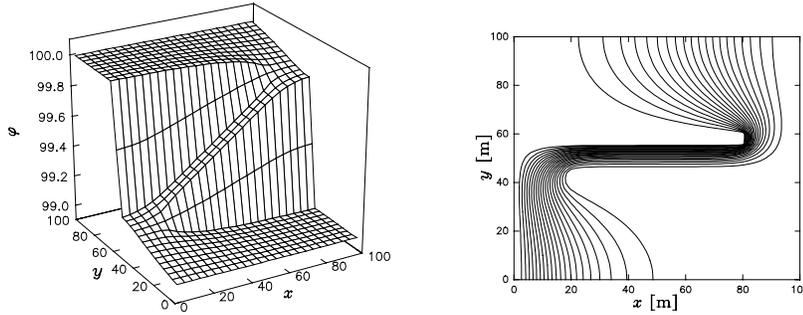


FIG. 4.2. Solution $\phi(x, y)$ and streamlines of the first example (above) in Mosé et al. [21].

TABLE 4.5
Iteration counts with single $V(1, 1, 1)$ -cycles.

Mesh size	$CG(\mathcal{S}_\phi)$	\mathcal{S}_ϕ^D	$\widetilde{\mathcal{S}}_{\mu;u}$	$\widetilde{\mathcal{S}}_{\mu;u,v}$	$\widetilde{\mathcal{S}}_\phi$	$CG(\mathcal{S}_\phi)$
50×50	795	342	6	4	11	667
100×100	1,675	600	6	4	11	1,709
200×200	3,049	1,027	6	4	11	4,486

TABLE 4.6
Iteration timings, [s] for the flow problem employing single $V(1, 1, 1)$ -cycles.

Mesh size	$CG(\mathcal{S}_\mu)$	\mathcal{S}_μ^D	$\widetilde{\mathcal{S}}_{\mu;u}$	$\widetilde{\mathcal{S}}_{\mu;u,v}$	$\widetilde{\mathcal{S}}_\phi$	$CG(\mathcal{S}_\phi)$
50×50	7.62	3.50	0.430	0.610	0.450	5.95
100×100	72.3	32.4	2.13	3.26	2.51	71.5
200×200	601	243	10.4	15.53	15.8	939

codes designed to solve the edge-based system. Finally, we note that on the 200×200 problem the approximate Schur complement preconditioners are approximately 16 times faster than the simple diagonal preconditioner.

4.3. A diffusive checkerboard. The final test of the preconditioners combines huge jumps in the diffusion coefficient with a spatially dependent grid containing high aspect ratio cells. Shown schematically in Figure 4.3, we consider a single checkerboard on the square domain $[0, 24] \times [0, 24]$ with $\Omega_1 = \{(0, 12) \times (0, 12)\} \cup \{(12, 24) \times (12, 24)\}$, $\Omega_2 = \{(0, 12) \times (0, 12)\} \cup \{(12, 24) \times (12, 24)\}$. We begin by introducing reflective boundary conditions

$$\mathbf{J} \cdot \mathbf{n} = 0 \quad \forall (x, y) \in \Gamma_R$$

along $\Gamma_R = \{x = 0, 0 < y < 24\} \cup \{0 < x < 24, y = 0\}$ and vacuum boundary conditions³

$$\frac{1}{4}\phi - \frac{1}{2}\mathbf{J} \cdot \mathbf{n} = 0 \quad \forall (x, y) \in \Gamma_V$$

along $\Gamma_V = \{x = 24, 0 < y < 24\} \cup \{0 < x < 24, y = 24\}$. In the constant coefficient case with a constant source this imposes a natural flow of neutrons along streamlines $y - x = \text{constant}$. Introducing a piecewise constant definition of the diffusion coefficient and the source

$$D(x, y) = \begin{cases} 1000 & \forall (x, y) \in \Omega_1, \\ 1 & \forall (x, y) \in \Omega_2, \end{cases} \quad Q(x, y) = \begin{cases} 1 & \forall (x, y) \in \Omega_1, \\ 0 & \forall (x, y) \in \Omega_2 \end{cases}$$

creates a significant perturbation of this idealized flow. Particularly, the relatively low diffusion coefficient in Ω_2 serves to channel the neutrons through a single point, $(12, 12)$, creating the large gradients observed in Figure 4.4. Thus, this problem introduces a mixed boundary condition and a significantly worse singularity than was present in the ground water flow problem (section 4.2). To further test the preconditioners and to accommodate internal layers near the interfaces of Ω_1 and Ω_2 we introduce an exponential mesh. Specifically, on the region $0 \leq x \leq 12$ we consider a grid spacing which is uniform in area such that

$$\Delta x = \frac{2}{L} \int_0^{12} e^{\gamma_x x} dx, \quad x_{i+1} = \frac{1}{\gamma_x} \log(e^{\gamma_x x} + \gamma_x \Delta x).$$

Similar definitions are employed for $12 \leq x \leq 24$ and for the grid in y . A typical mesh generated with $\gamma_x = \gamma_y = 2/10$ is displayed in Figure 4.5.

The iteration counts for this grid are given in Table 4.7, where column r indicates the maximum aspect ratio present. We first note that the iteration counts for $CG(\mathcal{S}_\mu)$ and $CG(\mathcal{S}_\phi)$ are comparable and extremely large (16,055 and 19,054, respectively, on the 96×96 mesh), giving a clear indication of the conditioning of the problem. The influence of \mathcal{S}_μ^D is quite impressive given its simplicity, achieving iteration counts which are only 25% higher than for the ground water flow problem. As expected the iteration counts for $\mathcal{S}_{\mu;u}$ are extremely poor, remaining consistent with the bound given in Theorem 3.2. Conversely, the results for both $\widetilde{\mathcal{S}}_{\mu;u,v}$ and $\widetilde{\mathcal{S}}_\phi$ remain excellent. We also note that a comparison of the iteration counts for $\widetilde{\mathcal{S}}_{\mu;u,v}$ in Table 4.7 with those for the constant coefficient Dirichlet problem in Table 4.3 with $r = 8$ (or $r = 1/8$) reveals a comparable performance on this significantly more difficult problem. The equivalent comparison for $\widetilde{\mathcal{S}}_\phi$ leads to the same conclusion.

³This mixed boundary condition is referred to as the vacuum boundary condition in reactor physics because it is derived using the P_1 (i.e., diffusion) approximation of the transport theory representation of the idealized material/vacuum interface. Physically it approximates the condition that no external source be applied to the boundary Γ_V while neutrons may exit from it.

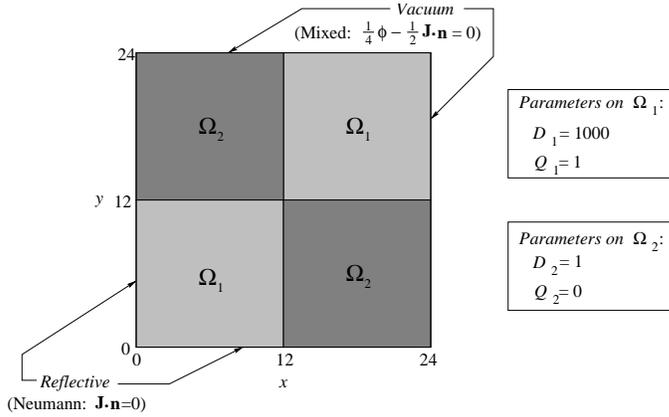


FIG. 4.3. A diffusive checkerboard configuration based on Dendy's [8] example.

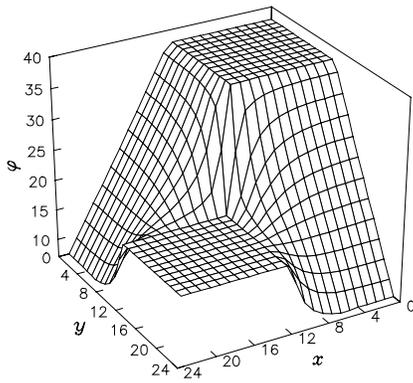


FIG. 4.4. The solution, $\phi(x,y)$ of the diffusive checkerboard computed on a 24×24 uniform mesh.

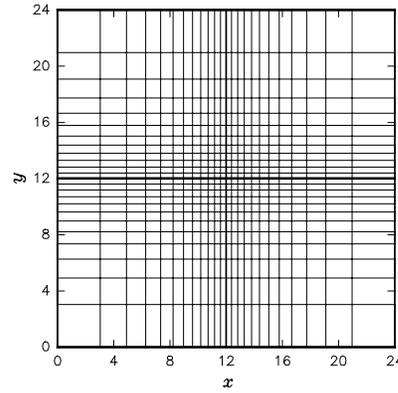


FIG. 4.5. Two-dimensional exponential grid with the parameters $\gamma_x = \gamma_y = 2/10$.

TABLE 4.7
Iteration counts for $\gamma_x = \gamma_y = 2/10$.

Mesh size	r	$CG(\mathcal{S}_\mu)$	\mathcal{S}_μ^D	$\widetilde{\mathcal{S}}_{\mu;u}$	$\widetilde{\mathcal{S}}_{\mu;u,v}$	$\widetilde{\mathcal{S}}_\phi$	$CG(\mathcal{S}_\phi)$
24×24	7.71	2,129	181	59	9	13	2,063
48×48	9.04	6,190	374	74	10	13	6,932
96×96	9.91	16,055	759	86	11	14	19,054

5. Conclusions. The families of nodal methods and mixed-hybrid finite element methods embody many desirable properties of the underlying PDE which contribute to their robustness and accuracy. Unfortunately, their very design makes the solution of the resulting equations awkward and costly, thwarting many potential users. Yet, inherent in this structure is a natural partitioning of the system which can be utilized in the development of preconditioners. The existence of such a partitioning suggests the investigation of various reduced systems (e.g., Schur complements) in

conjunction with suitably sparse approximations. In particular, we presented three such approximations, $\widetilde{\mathcal{S}}_{\mu;u}$, $\widetilde{\mathcal{S}}_{\mu;u,v}$, and $\widetilde{\mathcal{S}}_{\phi}$, and demonstrated that the preconditioning which resulted was optimal in the sense that a fixed number of iterations, independent of the mesh spacing, was required to reduce the residual by a fixed amount. These solvers are competitive with multilevel methods because the preconditioner is only approximately inverted with a single V- or W-multigrid cycle.

Unfortunately, $\widetilde{\mathcal{S}}_{\mu;u}$, which preconditions the most popular form of the discretization, is sensitive to high aspect ratio cells having $r > 1$. This shortcoming, an artifact of arbitrarily approximating A_{uu} as opposed to A_{vv} , was alleviated by the *two-step* preconditioner $\widetilde{\mathcal{S}}_{\mu;u,v}$. Conversely, an identical approximation is made in both coordinate directions during the development of $\widetilde{\mathcal{S}}_{\phi}$ and hence its preconditioning is insensitive to high aspect ratio cells. Moreover, the bound obtained in Theorem 3.4 suggests that in a more general setting the system $\mathcal{S}_{(\phi,\mu)}$ will be of significant practical interest. In all cases the preconditioners were found to be robust with respect to spatial variations in the diffusion coefficient.

Acknowledgments. The authors thank Dr. Joel Dendy for insightful observations and discussions that served as a catalyst for this research. The first author also thanks Dr. Jim Varah for many helpful discussions regarding the analysis of these solvers.

REFERENCES

- [1] M. B. ALLEN, R. E. EWING, AND P. LU, *Well-conditioned iterative schemes for mixed finite-element models of porous-media flows*, SIAM J. Sci. Statist. Comput., 13 (1992), pp. 794–814.
- [2] T. ARBOGAST AND Z. CHEN, *On the implementation of mixed methods as nonconforming methods for second-order elliptic problems*, Math. Comp., 64 (1995), pp. 943–972.
- [3] D. N. ARNOLD AND F. BREZZI, *Mixed and nonconforming finite element methods: Implementation, postprocessing and error estimates*, Math. Model. Numer. Anal., 19 (1985), pp. 7–32.
- [4] Y. Y. AZMY AND B. L. KIRK, *Performance modeling of parallel algorithms for solving neutron diffusion problems*, Nuclear Sci. Engrg., 120 (1995), pp. 1–17.
- [5] F. BREZZI AND M. FORTIN, *Mixed and Hybrid Finite Element Methods*, Springer Ser. Comput. Math. 15, Springer-Verlag, New York, 1991.
- [6] Z. CHEN AND D. KWAK, *The Analysis of Multigrid Algorithms for Nonconforming and Mixed Methods for Second Order Elliptic Problems*, IMA Preprint 1277, Institute for Mathematics and its Applications, Minneapolis, MN, 1994.
- [7] S. DEMKO, W. F. MOSS, AND P. W. SMITH, *Decay rates for inverses of band matrices*, Math. Comp., 43 (1984), pp. 491–499.
- [8] J. E. DENDY, *Black box multigrid*, J. Comput. Phys., 48 (1982), pp. 366–386.
- [9] J. DOUGLAS, R. DURÁN, AND P. PIETRA, *Numerical Approximation of Partial Differential Equations*, North-Holland, Amsterdam, 1987, Chapter: Formulation of Alternating-Direction Iterative Methods for Mixed Methods in Three Space, pp. 21–30.
- [10] R. E. EWING, Y. KUZNETSOV, R. D. LAZAROV, AND S. MALIASSO, *Substructure preconditioning for porous flow problems*, in Finite Element Modeling of Environmental Problems, John Wiley, New York, 1995, pp. 303–329.
- [11] H. FINNEMANN, F. BENNEWITZ, AND M. R. WAGNER, *Interface current techniques for multidimensional reactor calculations*, Atomkernenergie, 30 (1977), pp. 123–128.
- [12] H. D. FISCHER AND H. FINNEMANN, *The nodal integration method – A diverse solver for neutron diffusion problems*, Atomkernenergie–Kerntechnik, 39 (1981), pp. 229–236.
- [13] N. K. GUPTA, *Nodal methods for three-dimensional simulators*, Progr. Nuclear Energy, 7 (1981), pp. 127–149.
- [14] J. P. HENNART AND E. DEL VALLE, *Fast elliptic solvers using mixed-hybrid nodal finite elements*, Advances in Numerical PDE’s and Optimization, S. Gómez, J. P. Hennart, and R. A. Tapia, eds., SIAM, Philadelphia, PA, 1991, pp. 171–184.

- [15] J. P. HENNART AND E. DEL VALLE, *On the relationship between nodal schemes and mixed-hybrid finite elements*, Numer. Methods Partial Differential Equations, 9 (1993), pp. 411–430.
- [16] Y. A. KUZNETSOV AND S. Y. MALIASSOV, *Substructuring preconditioners for nonconforming finite element approximations of second order elliptic problems with anisotropy*, Russian J. Numer. Anal. Model, 10 (1995), pp. 511–533.
- [17] R. D. LAWRENCE, *Progress in nodal methods for the solution of the neutron diffusion and transport equations*, Progr. Nuclear Energy, 17 (1986), pp. 271–301.
- [18] R. D. LAWRENCE, *Three-dimensional nodal diffusion and transport methods for the analysis of fast-reactor critical experiments*, Progr. Nuclear Energy, 18 (1986), pp. 101–111.
- [19] R. D. LAWRENCE AND J. J. DORNING, *A nodal green's function method for multidimensional neutron diffusion calculations*, Nuclear Sci. Engrg., 76 (1980), pp. 218–231.
- [20] B. MONTAGNINI, P. SORAPERRA, C. TRENTAVIZI, M. SUMINI, AND D. M. ZARDINI, *A well-balanced coarse-mesh flux expansion method*, Ann. Nuclear Energy, 21 (1994), pp. 45–53.
- [21] R. MOSÉ, P. SIEGEL, P. ACKERER, AND G. CHAVENT, *Application of the mixed hybrid finite element approximation in a groundwater flow model: Luxury or necessity*, Water Resource Res., 30 (1994), pp. 3001–3012.
- [22] C. R. RAO AND S. K. MITRA, *Generalized Inverse of Matrices and its Applications*, John Wiley, New York, 1971.
- [23] T. RUSTEN AND R. WINTHER, *A preconditioned iterative method for saddlepoint problems*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 887–904.
- [24] J. W. SONG AND J. K. KIM, *An efficient nodal method for transient calculations in light water reactors*, Nuclear Tech., 103 (1993), pp. 157–167.
- [25] M. R. WAGNER, H. FINNEMANN, K. KOEBKE, AND H.-J. WINTER, *Validation of the nodal expansion method and the depletion program MEDIUM-2 by benchmark calculations and direct comparison with experiment*, Atomkernenergie, 30 (1977), pp. 129–135.
- [26] D. M. YOUNG AND K. C. JEA, *Generalized conjugate-gradient acceleration of nonsymmetrizable iterative methods*, Linear Algebra Appl., 34 (1980), pp. 159–194.